

**APPLICATION FOR  
UNITED STATES PATENT  
IN THE NAME OF**

**DAYLAN B. DARBY**

Assigned to

**INTEL CORPORATION**

for

**ENCODED ELECTRONIC MAIL**

Prepared by:

Pillsbury Winthrop LLP  
725 South Figueroa Street, Suite 2800  
Los Angeles, California 90017-5406  
Tel: (213) 488-7100  
Fax: (213) 629-1033

ATTORNEY DOCKET NO. 081674-0280173

Express Mail No. EM EL 724 027 404 US

081674-0280173

## ENCODED ELECTRONIC MAIL

### BACKGROUND

#### 1. Field of the Invention

Embodiments described herein are directed to encoded electronic mail that includes the files that represent the meta-data of an email, the files that represent the email data, and the processes that act on these files. The meta-data is combined with the email data within a single file, and encoding, by means of a header, is implemented to monitor the changes to the meta-data and the location of the email data within the file.

#### 2. Related Art

Current electronic mail systems utilize the files that represent the meta-data of electronic mail as well as the files that represent the electronic mail data. An email process receives the email data into a data-file and creates one or more meta-files to describe the email. Other email functions consult and/or modify the meta-files to process (e.g., parse, route, or forward) the email. This method is prone to email corruption. That is, if any of the meta-files are deleted inadvertently, the corresponding data-file becomes unusable. The converse is also true. Several causes of accidental deletion include a system crash, failed backup/restore, and administrative error.

In a typical operating system, email files are treated in a like manner as other files. The operating system does not have an inherent understanding of the relationship between files. For instance, in a conventional email system, a certain set of the email meta-files is considered the root or main files. These files contain references such as filename and file extensions to other meta-files and data-files. If a root is accidentally deleted, the other meta-files and data-files associated with it are orphaned such that they consume disk space without being referenced. If a non-root file is deleted, then the set of meta-files and data-files that represent that particular email message are incomplete.

Moreover, in a representative email system, the set of meta-files and data-files that represent a particular email message are all opened at once, thereby consuming operating system-limited file descriptors and file buffers. These opened files cause the operating system to allocate memory from the system load, thus reducing the amount of memory and file descriptors available to other non-email processes. To alleviate the problem, there is a need to minimize open file descriptors and therefore maximize system resources for non-email processes.

Encoded electronic mail is thus designed to minimize the use of the central processing unit by maximizing the use of system resources such as memory, disk space, file descriptors, locks, and the like. That is, the meta-data is combined with the actual email data within a single file, and encoding, via the addition of a header, is used to monitor the changes to the meta-data and the location of the actual email data within the file.

## **BRIEF DESCRIPTION OF THE DRAWINGS**

In order that all of the structural and functional features for attaining the objects for Encoded Electronic Mail may be readily understood, a detailed description of embodiments of the invention will be made with reference to the accompanying drawings, wherein like numerals designate corresponding parts in the several figures.

FIG. 1 is a block diagram illustrating the components of encoded electronic mail.

FIG. 2 is a flow chart illustrating the steps associated with receipt of an electronic mail.

FIG. 3 is a flow chart illustrating the steps associated with an electronic mail process.

FIG. 4 is a depiction of an email system in a data communications network.

## **DETAILED DESCRIPTION**

The following paragraphs describe encoded electronic mail according to an embodiment of the present invention. The components include the files that represent the meta-data of the electronic mail, the files that represent the electronic mail data, and the processes that act on these files. Meta-data files contain information such as, but not limited to, email sender, email receiver(s), email file size, email forwarding, and processing information.

In a typical system, when an email message is sent from one computer to another computer, the sending computer establishes a Simple Mail Transfer Protocol ("SMTP"), i.e., RFC821 session with the receiving computer. The receiving computer may be the ultimate destination or an intermediate destination. SMTP commands are generated by the sender SMTP and sent to the receiver SMTP. SMTP replies are sent from the receiver SMTP to the sender SMTP in response to the commands. The sending computer sends a "HELO . . ." and the receiving computer responds with an "OK . . ." The HELO command is used to identify the sender SMTP to the receiver SMTP wherein the argument field contains the host name of the sender SMTP. The receiver SMTP identifies itself to the sender SMTP in the greeting reply and in the response to the HELO command. This command and the OK reply to it confirm that the

sender SMTP and the receiver SMTP are in the initial state. That is, there is no transaction in progress.

The sending computer then sends a "MAIL FROM . . ." The MAIL command is used to initiate a mail transaction in which the mail data is delivered to one or more mailboxes. Provided that the receiving computer can accept the mail, it responds with an "OK . . ."

The sending computer sends one or more "RCPT TO . . ." indicating a recipient of the mail data; multiple recipients are specified by multiple use of the RCPT TO command. If the receiving computer can accept mail for those recipients, it responds to each with an "OK . . ." If not, it responds with a reply rejecting that recipient but not the entire mail transaction. The RCPT TO information is a variable part of the meta-data. An example of a RCPT TO is jane.doe@schwa.com. The RCPT TO may also be an alias, a group list, a forwarding account, or the like. In a conventional system, the date, MAIL FROM, and RCPT TOs are included in the meta-file. The receiving computer builds the meta-file during the conversation with the sending computer.

After all the RCPT TOs are sent, the sending computer sends "DATA" and the receiving computer responds with "OK". The sending computer then sends the RFC822 header and body, which both appear as DATA to the RFC821 session, and the receiving computer responds with "OK . . ." and the sending computer either sends the next email message or closes the SMTP.

In sum, there are three steps to SMTP mail transactions. The transaction commences with a MAIL command that gives the sender identification. A series of one or more RCPT commands follows giving the receiver information. Then, a DATA command gives the mail data. Finally, the end of a mail data indicator confirms the transaction.

The receiving computer then has a meta-file that contains the RFC821 instructions and at least a data file that contains the RFC822 header and body. Some mail systems split the data file into two files, whereby one contains the RFC822 header and the other contains the RFC822 body. The RFC822 header essentially is a human readable email routing instruction that is followed by a blank line and then the RFC822 body. Email attachments are always encoded in the RFC822 body.

The receiving computer then has to determine to whom to forward the email message. In a simple case, the meta-file, containing the RFC821 instructions, contains a single forwarding address, i.e., jane.doe@schwa.com, and the RFC822 header and body are either appended to the

local mailbox or forwarded onto the next computer. In a non-trivial case, the meta-file contains an alias, a group list, a forwarding account, etc. For instance, if a message is addressed to all@schwa.com instead of jane.doe@schwa.com, the receiving computers would replace the single "RCPT TO: all@schwa.com" with multiple names. Because it is an inefficient operation to insert names in a file, the email software would not alter the RFC822 header file for each insert. Rather, the email software would rewrite the rather short meta-file that contains the RFC821 instructions. The receiving computer would then either append the RFC822 header and body files to each local mailbox or become a sending computer and forward the email message to the next computer.

From the point of reception of an email message to the point of either relaying the message to another computer or storing the message in a local mailbox, the email software typically treats the RFC822 header and the RFC822 data as read-only. Inserting and deleting routing information in the RFC822 header on a continuous basis is an inefficient practice. As such, the insertion, deletion, and other manipulation functions are only performed in the meta-file. Only when the email message is ready to be stored locally or sent to another computer is the RFC822 header rewritten with the new information.

In one embodiment of the present invention, as illustrated in Figure 1, variable meta-data information **130** is combined with actual email data **120** within a single file **140**. A header is then used to monitor any changes to the meta-data **130** and the location of the actual email data **120** within the file. A format of the combined meta-file/data-file **140** could be, but is not limited to, a fixed sized header **110**, that contains links or indices to the information within the file **140**. This includes a link to the header size. The header **110** may be expressed in any common syntax including, but not limited to, XML, HTML, and numeric offsets. The header **110** links include a link to the start of the email data, a link to the email sender, a link to the email receiver(s), and the like. All the meta-data information **130** is similarly referenced, but not necessarily stored, in the header **110**.

Following the header **110** is the actual email data **120**. The data portion is of variable length. Once recorded, the data portion does not change in size. Following the header **110** and the actual email data **120** is the variable meta-data information **130**. Because email recipients are often aliased and forwarded and mailing lists are expanded, this information changes as the email

is processed within the email system. Rather than rewriting the meta-files, this section is rewritten or re-appended and the header links rewritten.

A possible format for the single file **140** that describes the internally stored email message is found in Table 1.

TABLE 1	
Received Time	32bit number of seconds since Epoch that this message was received
HELO Index	32bit number of bytes from start of file where HELO information is stored
MAIL FROM Index	32bit number of bytes from start of file where MAIL FROM information is stored
RCPT TO Index	32bit number of bytes from start of file where RCPT TO information is stored
RCPT TO Count	32bit number of RCPT TO strings
RFC822 Header Index	32bit number of bytes from start of file where RFC822 header information is stored
RFC822 Body Index	32bit number of bytes from start of file where RFC822 body information is stored
RFC822 Header	A NULL terminated character string containing the received RFC822 header
RFC822 Body	A NULL terminated character string containing the received RFC822 body
HELO Information	A NULL terminated character string containing the name of the sending computer
MAIL FROM Information	A NULL terminated character string containing the name of the sending user
RCPT TO Information	A set of NULL terminated character strings containing the name(s) of the recipient(s)

Because the RCPT TO information is variable, it is last in the chain such that an append/overwrite file operation can be used to modify the RCPT TO list. The same type of email processing described above would then be performed using this format. That is, as shown in Figure 2, an SMTP email connection is accepted. This is illustrated in step **210**. As shown in step **220**, email data **120** is then received. This email data **120** is then recorded in the data section of the single file **140**. This is shown in step **230**. As illustrated in step **240**, the variable meta-data information **130** is then computed and recorded. Next, header **110** links are computed and recorded. This is demonstrated in step **250**. Finally, as shown in step **260**, the email file **140** is passed on for processing.

As shown in Figure 3, regarding email processing, the combined single file **140** would be opened, as illustrated by step **310**. Based on aliases, forwarding rules, group lists, etc., the

variable meta-data information **130** is then recomputed, as shown by step **320**. Delivery, or possibly retransmission, of the message is then attempted. This is illustrated in step **330**. Next, as shown in step **340**, the variable meta-data information **130** is rewritten. Based on any adjustments, header **110** links are then changed in order to monitor the modified information.

5 This is illustrated in step **350**. Finally, the single file **140** is closed, as illustrated in step **360**.

Figure 4 depicts the overall electronic mail system in operation. That is, as previously described, a sending computer **402** establishes an SMTP connection with one or more receiving computers **404**. Transmission of the electronic mail message occurs over a data communication network **406** to which the sending computer **402** and the one or more receiving computers **404**  
10 are linked. The data communication network **406** may include the Internet, an Intranet, or any combination of public and private data communication networks.

While the above description refers to particular embodiments of the present invention, it will be understood to those of ordinary skill in the art that modifications may be made without departing from the spirit thereof. The accompanying claims are intended to cover any such  
15 modifications as would fall within the true scope and spirit of the present invention.

The presently disclosed embodiments are therefore to be considered in all respects as illustrative and not restrictive; the scope of the invention being indicated by the appended claims, rather than the foregoing description. All changes that come within the meaning and range of equivalency of the claims are therefore intended to be embraced therein.